

CLAIMS

What is claimed is:

1. A method of buffering at least one data unit received at a node on a communications network, the at least one data unit being associated with one or more channels in the network, the network node including at least one input port communicably coupleable to at least one output port, comprising the steps of:

providing a memory at each output port of the network node, the memory comprising at least one linear time-indexed array having a plurality of locations for buffering the at least one data unit;

in the event the at least one data unit is associated with a single channel in the network, storing the at least one data unit in a respective location of the time-indexed array; and

in the event the at least one data unit is associated with a plurality of channels in the network, dividing the memory into a plurality of linear time-indexed arrays, each time-indexed array corresponding to a respective channel, and storing the at least one data unit in a respective location of the corresponding time-indexed array.

2. The method of claim 1 wherein the network node has a predetermined total bandwidth, and the providing step includes providing a memory within the node, the memory having a size sufficient to support the total bandwidth of the node.

3. The method of claim 1 wherein the dividing step includes dividing the memory into a plurality of arrays, each array corresponding to a respective channel, the respective channels conforming to predetermined bandwidth requirements.

4. The method of claim 3 wherein the dividing step includes dividing the memory into a plurality of arrays, each array having a size proportional to a fractional amount of a predetermined total bandwidth of the node.

5. The method of claim 1 wherein the at least one data unit has an associated timestamp value, and the first storing step includes storing the at least one data unit in a respective location of the time-based array based on the associated timestamp value.

6. The method of claim 5 wherein the first storing step includes storing the at least one data unit in a respective location of the time-based array in accordance with a weighted-fair queuing algorithm.

7. The method of claim 1 wherein the at least one data unit has an associated timestamp value, and the second storing step includes storing the at least one data unit in a respective location of the corresponding array based on the associated timestamp value.

8. The method of claim 7 wherein the second storing step includes storing the at least one data unit in a respective location of the corresponding array in accordance with a weighted-fair queuing algorithm.

9. A method of scheduling the transmission of at least one data unit from a node on a communications network, the network

node including at least one input port communicably coupleable to at least one output port, comprising the steps of:

providing a first memory at each output port of the network node, the first memory comprising at least one linear time-indexed array having a plurality of locations for buffering the at least one data unit;

receiving at least one first data unit at the network node, the at least one first data unit having an associated timestamp value;

inserting the first data unit into a respective location of the time-based array based on the associated timestamp value;

partitioning a binary value of the timestamp associated with the first data unit into a plurality of sub-fields, each sub-field comprising one or more bits, and using the plurality of sub-fields to generate a corresponding plurality of acceleration bit-strings for use in identifying the first data unit in the time-based array having a lowest associated timestamp value;

extracting the identified first data unit having the lowest associated timestamp value from the time-based array; and

designating the extracted first data unit as a next data unit to be transmitted over the network.

10. The method of claim 9 wherein the network node includes at least one second memory and the partitioning step includes using the sub-fields of bits to index respective locations in the at least one second memory and asserting values at the respective memory locations to generate the plurality of acceleration bit-strings.

11. The method of claim 10 further including the step of priority encoding each acceleration bit-string to obtain a corresponding priority-encoded acceleration bit-string.

5 12. The method of claim 11 wherein the priority encoding step employs "low-wins" priority encoding.

10 13. The method of claim 11 further including the step of employing one or more of the priority-encoded acceleration bit-strings to index the time-based array to identify the first data unit in the array having the lowest associated timestamp value.

14. The method of claim 9 wherein each location of the time-based array corresponds to a respective timestamp value within a first time window ranging from $t=0$ to $t=T_w$, and further including the step of in the event the timestamp value associated with the next data unit to be transmitted over the network is greater than or equal to $T_w/2$, shifting the first time window forward in time by $T_w/2$ to obtain a next time window ranging from $t=T_w/2$ to $t=3T_w/2$.

15. The method of claim 14 wherein the receiving step includes receiving at least one first data unit at the network node, the at least one first data unit having an associated timestamp value within a range limited to $T_w/2$.

16. A system for scheduling the transmission of at least one data unit from a node on a communications network, the node including at least one input port and at least one output port, the input port being communicably coupleable to the output port, comprising:

1056609-012402
201211-6095500
20
a first memory disposed at each output port of the network node, the first memory comprising a linear time-indexed array having a plurality of locations configured to buffer at least one first data unit, each first data unit having an associated timestamp value; and

5 a controller configured to insert the at least one first data unit into a respective location of the time-based array based on the associated timestamp value, partition a binary value of the timestamp associated with the first data unit into a plurality of sub-fields, each sub-field comprising one or more bits, use the plurality of sub-fields to generate a corresponding plurality of acceleration bit-strings for use in identifying the first data unit in the time-based array having a lowest associated timestamp value, extract the identified first data unit having the lowest associated timestamp value from the time-based array, and designate the extracted first data unit as a next data unit to be transmitted over the network.

17. The system of claim 16 wherein each location of the time-based array corresponds to a respective timestamp value within a first time window ranging from $t=0$ to $t=T_w$, and the controller is further configured to, in the event the timestamp value associated with the next data unit to be transmitted over the network is greater than or equal to $T_w/2$, shift the first time window forward in time by $T_w/2$ to obtain a next time window ranging from $t=T_w/2$ to $t=3T_w/2$.

18. The system of claim 16 wherein the network node has a predetermined total bandwidth and the time-based array has a size sufficient to support the total bandwidth of the node.

19. The system of claim 16 wherein the first memory comprises a plurality of linear time-indexed arrays, each array corresponding to a respective channel in the network.

5 20. The system of claim 19 wherein each array has a size proportional to a fractional amount of a predetermined total bandwidth of the network node.